

SYSTEM AND METHOD FOR DELIVERY OF SHORT-TIME DURATION VIDEO SEGMENTS

5

RELATED APPLICATIONS

The present application claims priority from U.S. Provisional Patent Application Serial No. 60/141,533, filed June 28, 1999, inventors Sadik Bayrakeri and Donald F. Gordon, and entitled "SERVER-CENTRIC SHORT-TIME DURATION VIDEO SEGMENT DELIVERY PROCESS IN AN INTERACTIVE TELEVISION ENVIRONMENT." The present application also claims priority from U.S. Provisional Patent Application Serial No. 60/178,100, filed January 26, 2000, inventors Sadik Bayrakeri and Donald F. Gordon, and entitled "BANDWIDTH MANAGEMENT TECHNIQUES FOR DELIVERY OF INTERACTIVE PROGRAM GUIDE."

The present application is a continuation-in-part of co-pending U.S. Patent Application Serial No. 09/524,854, filed March 14, 2000, inventors Sadik Bayrakeri, Donald F. Gordon, Edward A. Ludvig, Eugene Gershtein, Jeremy S. Edmonds, John P. Comito, and Alfred Li, and entitled "BANDWIDTH MANAGEMENT TECHNIQUES FOR DELIVERY OF INTERACTIVE PROGRAM GUIDE." The present application is also a continuation-in-part of co-pending U.S. Patent Application Serial No. 09/538,562, filed March 29, 2000, inventors Donald F. Gordon, Sadik Bayrakeri, Edward A. Ludvig, Eugene Gershtein, Jeremy S. Edmonds, John P. Comito, and Alfred Li, and entitled "DEMAND-CAST SYSTEM AND BANDWIDTH MANAGEMENT FOR DELIVERY OF INTERACTIVE PROGRAMMING." The present application is also a continuation-in-part of co-pending U.S. Patent Application Serial No. 09/539,228, filed March 30, 2000, inventors Donald F. Gordon, Sadik Bayrakeri, Edward A. Ludvig, Eugene Gershtein, Jeremy S. Edmonds, John P. Comito, and Alfred Li, and entitled "MESSAGING PROTOCOL FOR DEMAND-CAST SYSTEM AND BANDWIDTH MANAGEMENT."

BACKGROUND OF THE INVENTION

1. Field of the Invention

The invention relates to communications systems in general. More specifically, the invention relates to video communications systems.

5

2. Description of the Background Art

Over the past few years, the television industry has seen a transformation in a variety of techniques by which its programming is distributed to consumers. Cable television systems are doubling or even tripling system bandwidth with the migration to hybrid fiber coax (HFC) cable plant. Customers unwilling to subscribe to local cable systems have switched in high numbers to direct broadcast satellite (DBS) systems. And, a variety of other approaches have been attempted focusing primarily on high bandwidth digital technologies, intelligent two way set top terminals, or other methods of trying to offer service differentiated from standard cable and over the air broadcast systems.

With this increase in bandwidth, the number of programming choices has also increased. Leveraging off the availability of more intelligent set top terminals, several companies such as Starsight Telecast Inc. and TV Guide, Inc. have developed elaborate systems for providing an interactive listing of a vast array of channel offerings, expanded textual information about individual programs, the ability to look forward to plan television viewing as much as several weeks in advance, and the option of automatically programming a VCR to record a future broadcast of a television program.

Unfortunately, conventional electronic program guides have several drawbacks. They tend to require a significant amount of memory, some of them needing upwards of one megabyte of memory at the set top terminal (STT). They are very slow to acquire their current database of programming information when they are turned on for the first time or are subsequently restarted (e.g., a large database may be downloaded to a STT using only a vertical blanking interval (VBI) data insertion technique). Such slow database acquisition may result in out of date database information or, in the case of a pay per view (PPV) or video on demand (VOD) system, limited scheduling flexibility for the information provider.

Figure 3 depicts a first pull topology for demand-casting video segments in accordance with an embodiment of the present invention.

Figure 4 is a flow chart showing a second pull method for demand-casting video segments in accordance with an embodiment of the present invention.

5 Figure 5 depicts a second pull topology for demand-casting video segments in accordance with an embodiment of the present invention.

Figure 6 depicts a two-way system for efficient delivery of demand-cast video segments in accordance with an embodiment of the present invention.

10 Figure 7 illustrates an example of an interactive program guide (IPG) screen in accordance with an embodiment of the present invention.

Figure 8 depicts an embodiment for the content of the demand-cast index table.

Figure 9 depicts one embodiment for the contents of the messages sent from the terminal to the SM .

15 Figure 10 depicts one embodiment for the contents of the messages sent from the SM to the TSG.

Figure 11 depicts one embodiment for the contents of the acknowledgement messages sent by the TSG back to the SM.

20 Figure 12 depicts an example showing status of active demand-cast streams in a multiplexed transport stream generated by a TSG.

FIG. 13 depicts a matrix of program guide data configured to present a different video for each PID.

DESCRIPTION OF THE SPECIFIC EMBODIMENTS

25 I. ILLUSTRATIVE COMMUNICATIONS NETWORK

Figure 1 depicts an illustrative communications network 100 for distributing video segments to a plurality of terminals in accordance with an embodiment

of the present invention. The illustrative network 100 comprises a cable distribution network, but other types of distribution networks may also be used within the spirit and scope of the present invention.

5 The illustrative network 100 includes one or more head-ends 102, one or more centers for local neighborhood equipment 104, a plurality of distribution nodes 106, and a plurality of subscriber stations 108. The local neighborhood equipment (LNE) 104 may be located, for example, at remote hubs of a cable distribution network. The end-user terminals 108 may comprise, for example, interactive set-top terminals (STT) or other devices with similar interactive functionalities.

10 II. EXAMPLE METHODS AND TOPOLOGIES

In the present application, the demand-cast term is used to refer to the process of managing and delivering content to one or more users depending on user demand for the content. Figures 2-7 depicts various methods and topologies for demand-casting video segments in accordance with embodiments of the present invention. The
15 various methods/topologies are given for purposes of edification and are not meant to limit the scope of the present invention.

Sub A' >
Figure 2 is a flow chart showing a first pull method 200 for demand-casting video segments in accordance with an embodiment of the present invention. The first pull method 200 involves pointcasting and provides a requesting viewer with access
20 to a video segment starting at its beginning (or any other selectable point within the video segment). As described below, the method 400 includes three steps.

In a first step 202, a request for a video segment is received from a STT 108. The request is transmitted upstream from the STT 108 to the HE 102 or LNE 104 by way of the communications network 100. The upstream transmission may be done via an
25 out-of-band network. Alternatively, the upstream transmission may be done via an in-band network.

In a second step 204, bandwidth to pointcast the requested video segment is allocated in the distribution system for that purpose. For example, as described in more detail below, a bandwidth manager (BWM) within a head-end 102 and/or local
30 neighborhood equipment 104 may allocate within the in-band network the necessary

bandwidth to pointcast the requested video segment to the requesting STT 108. Such allocation is performed if sufficient system resources are available to establish such a session.

5 In a third step 206, the requested video segment is pointcast to the requesting set-top terminal (STT) 108. The pointcast need only be received by the requesting STT 108 and does not need to be received by other STTs 108. The pointcast is sent downstream from the head-end 102 or local neighborhood equipment 104 to the requesting STT 108. The pointcast is performed within the allocated in-band bandwidth.

10 Figure 3 depicts a first pull topology for demand-casting video segments in accordance with an embodiment of the present invention. The topology 300 relates to the pull method 200 of Fig. 2. As shown in Fig. 3, the request is transmitted upstream from the requesting STT 108 to the HE 102 or LNE 104 via illustrative communications network 100. Subsequently, the requested video segment is pointcast downstream from the HE 102 or LNE 104 to the requesting STT 108 via the network 100.

15 Figure 4 is a flow chart showing a second pull method 400 for demand-casting video segments in accordance with an embodiment of the present invention. The method involves narrowcasting and also provides a requesting viewer with access to a video segment starting at its beginning (or any other selectable point within the video segment). As described below, the method 400 includes three steps.

20 In a first step 402, a request for a video segment is received from a requesting STT 502. The request is transmitted upstream from the requesting STT 502 to the HE 102 or LNE 104 by way of the communications network 100. The upstream transmission may be done via an out-of-band network. Alternatively, the upstream transmission may be done via an in-band network.

25 In a second step 404, bandwidth to narrowcast the requested video segment is allocated in the distribution system for that purpose. For example, as described below in relation to Figs. 10 through 13, a bandwidth manager (BWM) within a head-end 102 and/or local neighborhood equipment 104 may allocate within the in-band network the necessary bandwidth to narrowcast the requested video segment to a group
30 504 of terminals which includes the requesting STT 502. Such allocation is performed if sufficient system resources are available to establish such a session. The group 504 may

include terminals 108 in one geographic area or terminals 108 dispersed among different geographic areas but linked, for example, via a network group address.

5 In a third step 406, the requested video segment is narrowcast to the group 504 of terminals 108. The narrowcast need only be received by terminals 108 within the group 504 and does not need to be received by other STTs 108. The narrowcast is sent downstream from the head-end 102 or local neighborhood equipment 104 to the group 504 of terminals 108. The narrowcast is performed within the allocated in-band bandwidth.

10 Figure 5 depicts a second pull topology 500 for demand-casting video segments in accordance with an embodiment of the present invention. The topology 500 relates to the pull method 400 of Fig. 4. As shown in Fig. 5, the request is transmitted upstream from the requesting STT 502 to the HE 102 or LNE 104 via illustrative communications network 100. Subsequently, the requested video segment is narrowcast downstream from the HE 102 or LNE 104 to the group 504 which includes the requesting
15 STT 502 via the network 100.

III. DEMAND-CAST SYSTEM

1. Demand-cast Overview

20 In accordance with an embodiment of the present invention, video segments may be delivered using a demand-cast system. The demand-cast system may be, for example, a two-way system requiring communication between STT users on the cable network and the head-end via a back-channel.

25 In accordance with one embodiment of the present invention, such video segments may include multiple previews and advertisements incorporated into an interactive program guide (IPG). For example, a preview of a broadcast program may be made available via the IPG such that the user may select and play the preview. The following discusses demand-casting video segments incorporated into an IPG.

30 The demand-cast video segments may be inserted in multiplexed transport stream for temporary broadcast based on access demand generated by STT users on the cable network. When a request for a video segment is made by a particular STT, the STT requests from the head-end that the corresponding stream be inserted in the multiplexed

transport stream. The head-end then inserts the requested stream into the multiplexed transport stream.

When a STT requests that a new demand-cast video segment be inserted into the multiplexed transport stream, if there is no slot available in the multiplexed transport stream, the head-end refuses to insert a newly requested video segment resulting in a blockage. All statistical systems are susceptible to blockage if loaded with too many users or in the case of rare chaotic episodes.

2. System Description

Figure 6 depicts a two-way system 600 for efficient delivery of demand-cast video sequences in accordance with an embodiment of the present invention. The system 600 includes a session manager (SM) 602 and a transport stream generator (TSG) 604.

Both the SM 602 and the TSG 604 may be, for example, co-located at a distribution center. The distribution center may comprise, for example, a headend 102 in the illustrative distribution system 100. Alternatively, the SM 602 and the TSG 604 may be at different locations. For example, the SM 602 may be at a headend 102, and the TSG 604 may be at local neighborhood equipment 104 in the illustrative distribution system 100.

Both the SM 602 and the TSG 604 are coupled to a plurality of terminals 606 via a distribution network. The distribution network may comprise, for example, the cable distribution network 100 illustrated in Fig. 1. In that example, the terminals 606 would comprise set-top terminals 108 or the equivalent functionality integrated into a computer system or advanced television. Alternatively, for example, the distribution network may comprise a satellite communications system or another type of communications system (telephonic, wireless, etc.).

In one embodiment, the terminal 606 receives in-band communications from the TSG 604 and sends out-of-band (OOB) communications to the SM 602. In an alternative embodiment, the communications to the SM 602 may comprise upstream in-band communications.

The session manager (SM) 602 may comprise, in one embodiment, a computer system residing at a cable headend 102. The computer system may comprise, for example, a computer server running a version of the UNIX (or alternatively Windows) operating system. The computer system may receive out-of-band communications from the terminals 606 by way of a connection to the network controller. For example, the connection may comprise an Ethernet connection, and the network controller may comprise one from General Instruments (now part of Motorola) or another supplier. The computer system also communicates with and controls the transport stream generator 604 by way of a network connection such as an Ethernet connection.

The SM 602 manages delivery of the IPG to terminals 606 on multiple cable nodes each served by a separate multiplexed transport stream generated at a TSG 604. The SM 602 also monitors demand-cast stream usage by the terminals 606. It tracks IPG demand-cast streams that are acquired by at least one terminal 606 by maintaining a dynamic list of terminals 606 using each stream. This is done for each multiplexed transport stream managed by the SM 602. The SM 602 also accepts messages from terminals 606 indicating that they have acquired, released, or requested demand-cast streams. A terminal 606 that has acquired a demand-cast stream is registered to the stream, and a terminal 606 that has released a demand-cast stream is removed from the registry for the stream. The SM 602 informs the corresponding TSG 604 once there is no longer any terminals 606 registered to a particular demand-cast stream. It also informs the TSG 604 when a terminal 606 requests a demand-cast stream. In one embodiment, the SM 602 may time-out acquisition of a stream by any terminal 606 if the terminal 606 has not released the stream within a period of time (for example, a few minutes). The time-out may be implemented by removing the particular terminal 606 from the registry for the stream.

The transport stream generator (TSG) 604 may comprise, in one embodiment, a computer system residing at a cable headend 102. The computer system may comprise, for example, a computer server running a version of the Windows (or alternatively UNIX) operating system. Alternatively, the TSG 604 may be located separately from the SM 602, for example, at local neighborhood equipment 104. Each TSG 604 is coupled to a SM 602, for example, via an Ethernet network. The TSG 604

may generate one or more multiplexed transport stream, each broadcast to a respective node in the distribution system.

5 In one embodiment, the multiplexed transport stream comprises a MPEG transport stream. In this case, the TSG 604 may communicate with the terminals 606 by way of tables in the private section of the MPEG transport stream. Such a table may include a list of available demand-cast streams, along with the address of the source TSG 604 and information to identify the particular multiplexed transport stream to which the table belongs.

10 The TSG 604 manages each multiplexed transport stream which it generates. The TSG 604 receives information from the SM 602 indicating whether each demand-cast stream being served is currently being acquired by any terminal 606 or not. In other words, the TSG 604 is informed by the SM 602 when there is no longer any terminals 606 acquiring a demand-cast stream.

15 In one embodiment, the TSG 604 maintains a status for each variable demand-cast stream being served. The status is adjusted upon receipt by the TSG 604 of certain messages from the SM 602. The basic states for the status comprise an "acquired" state which denotes that the demand-cast stream is in use by one or more terminals 606, and a "released" state which denotes that the demand-cast stream is not in use by any terminal 606. The TSG 604 keeps serving "acquired" demand-cast streams by
20 multiplexing them into appropriate transport streams and replaces "released" demand-cast streams with new demand-cast streams upon receipt of a request message from the SM 602. In a preferred embodiment, the TSG 604 also keeps track of the order in which the streams are released, so that the oldest released stream may be used as the preferred candidate for replacement.

25 If all the demand-cast streams in a particular multiplexed transport stream are "acquired," then a new stream cannot be inserted, and so the TSG 604 refuses the request. In such a case, a message indicating such a refusal may be sent to the SM 602 and/or the requesting terminal 606.

30 In one embodiment, the terminal 606 comprises a set-top terminal (STT) for use by a service subscriber. The STT may comprise an embedded system which includes a tuner, a demultiplexer, and a decoder. The STT drives the subscriber's display

unit or TV set, and it may be connected to the TSG 604 by way of a RF feed from a cable distribution network. The IPG content may be received from a particular multiplexed transport stream on a specific QAM carrier signal. In one embodiment, the multiplexed transport stream may comprise an ensemble of elementary MPEG video streams, each
5 representing a portion of the guide.

In one embodiment, the terminal 606 includes client application software which is resident at the terminal 606. The client application is responsible for presenting the video segments to the subscriber. The client application demultiplexes and decodes video segments requested by the user. If a requested segment is already being received
10 via the multiplexed transport stream, then the Client application acquires the stream immediately and sends a message to the SM 602 that it has acquired the stream. If the requested segment is not in the multiplexed transport stream, then the client application sends a request message to the SM 602. Subsequently, the client application acquires the stream once it is received. In addition, when a stream is no longer being acquired, the
15 client application sends a release message to the SM 602. In one embodiment, if there is no remote control or other activity by the user for a period of time (for example, a few minutes), then the client application times-out the acquisition. The time-out may be accomplished, for example, by sending a release message to the SM 602 and acquiring a broadcast stream instead.

3. Description Per Major Module

The demand-cast system may include of three major subsystems: the set top terminal (STT); the session manager (SM); and the transport stream generator (TSG.)
25 A description of each subsystem follows.

A. STT (Set-Top Terminal)

The set top terminal may be the end-user or cable service subscriber
30 tuner/demultiplexer/decoder and embedded system. The STT may be an apparatus similar to the General Instruments DCT-2000. It is connected to the cable operator RF feed. It drives the subscribers display unit or TV set. The video segments may be received via a multiplexed transport stream (or multiplex, for brevity) located on a

specific QAM carrier. The multiplex contains an ensemble of elementary MPEG video streams each representing portions of the guide. Some of these streams are guide grid segments. The STT receives messages from the head-end via tables in the private section of the transport stream (in-band messaging.) The STT sends messages to the head-end via the out-of-band return path.

The client application may be a set top terminal resident program responsible for presenting an Interactive Program Guide including video segments to the subscriber. The client application demultiplexes and decodes video segments requested by the user. If a particular video segment is in the multiplexed transport stream, the STT acquires the stream immediately and informs the SM that it has requested it. If the video segment is not in the multiplex, the STT also sends a message to the SM that it has requested it. Then the STT acquires the stream once it's in the multiplex. When the STT no longer is acquiring the stream, it informs the SM that the stream has been released.

B. SM (Session Manager)

The session manager may be a computer residing at the cable head-end. The SM may be a SPARC Station running Solaris. It may be connected via Ethernet to the server side of a General Instruments network controller (NC) and may include a receiver for OB return path messages originating from STTs. The SM can handle STTs on multiple cable nodes each served by a separate multiplex. The SM communicates and controls the TSGs via Ethernet. The TSGs generate the transport streams.

The SM manages one or multiple cable networks and monitors demand-cast stream usage. It tracks demand-cast streams that are acquired by at least one STT maintaining a dynamic list of STTs that are using them. This is done for each multiplex managed by the SM. The SM accepts messages from STTs indicating that they have requested or released demand-cast streams. A STT that has acquired a demand-cast stream is registered to the stream and a STT that has released a demand-cast stream is removed from the streams registry. The SM informs the TSG once there are no longer any STTs on a particular demand-cast stream. It also informs the TSG when a STT requests a demand-cast stream.

C. TSG (Transport Stream Generator)

The transport stream generator is a computer residing at the cable head-end. Currently, the TSG is a PCI WinNT system. It is connected via Ethernet to the SM
5 controlling it. The TSG produces one or more transport streams each broadcast to their respective nodes. The TSG communicates with the STTs by way of tables in the private section of the transport streams. The table contains the list of available demand-cast streams along with the IP address of the source TSG (its IP address) and the channel number of the multiplex.

10

The TSG manages each multiplexed transport stream that it generates. It receives information from the SM on each demand-cast stream indicating whether the stream is currently acquired by any STT or released by all STTs. The TSG is informed by the SM when there is no longer any STT on a demand-cast stream. The TSG is
15 informed by the SM when a STT requests a demand-cast stream. The TSG maintains status for all the demand-cast streams in each multiplex. It adjusts the status for each stream for which it receives a message from the SM. The basic status is 'acquired' for streams in use by one or more STTs or 'released' for streams that are not in use by any STT. The TSG keeps 'acquired' streams in its multiplexes and replaces 'released'
20 streams with new demand-cast streams upon request. It also keeps track of which are the few oldest 'released' stream. The best candidate for replacement is always the oldest 'released' stream. If all the demand-cast streams in the multiplexes are 'acquired' then a new stream can not be inserted and the TSG refuses the request.

IV. EXAMPLE OF INTERACTIVE PROGRAM GUIDE

25 In accordance with an embodiment of the present invention, a video segment (and associated audio) may be delivered as part of an interactive program guide (IPG). An example of such an IPG is described below. The example is described for purposes of illustration and is not meant to limit the scope of the present invention.

30 *Sub A2* Figure 7 depicts an example of an IPG screen in accordance with an embodiment of the present invention. The video segment 700 of Figure 7 comprises a first 705A, second 705B and third 705C time slot objects, a plurality of channel content objects 710-1 through 710-8, a pair of channel indicator icons 741A, 741B, a video barker

A2
cancel.

5 720 (and associated audio barker), a cable system or provider logo 715, a program description region 750, a day of the week identification object 731, a time of day object 739, a next time slot icon 734, a temporal increment/decrement object 732, a "favorites" filter object 735, a "movies" filter object 736, a "kids" (i.e., juvenile) programming filter icon 737, a "sports" programming filter object 738 and a VOD programming icon 733. It should be noted that the day of the week object 731 and next time slot icon 734 may comprise independent objects (as depicted in Figure 7) or may be considered together as parts of a combined object.

10 In a system, illustratively, comprising 60 channels of information, the channels are displayed in 8-channel groups having associated with them three hour time slots. In this organization, it is necessary to provide 10 video PIDs to carry the present-time channel/time/title information, one or more audio PID to carry the audio barker and/or one or more data PID (or other data transport method) to carry the program description data, overlay data and the like. To fully broadcast interactive program

15 information up to 24 hours in advance, it is necessary to provide 160 ($10 \times 24 / 1.5$) video PIDS, along with one or more audio and, optionally, one or more data PIDs. The amount of time provided for in broadcast video PIDs for the given channel groups comprises the time depth of the program guide, while the number of channels available through the guide (compared to the number of channels in the system) provides the channel depth of

20 the program guide. In a system providing only half of the available channels via broadcast video PIDs, the channel depth is said to be 50%. In a system providing 12 hours of time slot "look-ahead," the time depth is said to be 12 hours. In a system providing 16 hours of time slot "look-ahead" and 4 hours of time slot "look-back," the time depth is said to be $+16/-4$ hours.

25 The video streams representing the IPG may be carried in a single transport stream or multiple transport streams, within the form of a single or multi-programs. A user desiring to view the next 1.5 hour time interval (e.g., 9:30 – 11:00) may activate a "scroll right" object (or move the joystick to the right when a program within program grid 702 occupies the final displayed time interval). Such activation results in

30 the controller of the STT noting that a new time interval is desired. The video stream corresponding to the new time interval is then decoded and displayed. If the corresponding video stream is within the same transport stream (i.e., a new PID), then the

stream is immediately decoded and presented. If the corresponding video stream is within a different transport stream, then the related transport stream is extracted from the broadcast stream and the related video stream is decoded and presented. If the corresponding transport stream is within a different broadcast stream, then the related
5 broadcast stream is tuned, the corresponding transport stream is extracted, and the desired video stream is decoded and presented.

A user interaction resulting in a prior time interval or a different set of channels results in the retrieval and presentation of a related video stream. If the related video stream is not part of the broadcast video streams, then a pointcast session, for
10 example, may be initiated as described above in relation to Figs. 2 and 3. For this purpose, the STT sends a request to the head end via the back channel requesting a particular stream. The head end then processes the request, retrieves the related stream from the information server, incorporates the stream within a transport stream as a video
15 PID (preferably, the transport stream currently being tuned/selected by the STT) and informs the STT which PID should be received, and from which transport stream it should be demultiplexed. The STT then retrieves the related video PID. In the case of the video PID being within a different transport stream, the STT first demultiplexes the corresponding transport stream (possibly tuning a different QAM stream within the forward channel).

20 Normally, upon completion of the viewing of the desired stream, the STT indicates to the head end that it no longer needs the stream, whereupon the head end tears down the pointcast session. The viewer is then returned to the broadcast stream from which the pointcast session was launched.

Frequently accessed video segments, such as those in the current time slot
25 and near look-ahead time slots and perhaps prime-time slots, may be broadcast constantly so as to remain accessible with low latency. Less frequently accessed far look-ahead pages may be pull demand-cast.

In accordance with an embodiment of the present invention, the IPG may contain video objects to provide access to video segments (and associated audio
30 segments). For example, the video barker 720 may be configured to be such a video object.

The user interface action sequence to initiate a request for the video segment may occur in various different ways. For instance, the video object may be selected using a select button on a remote when a cursor is on the video object .
Alternatively, other selection methods may be used, and additional functionality may be provided via pop-up menus and other user interface means to increase the user-friendliness of the terminal system. For instance, a separate remote button may be used to provide a pop-up menu with a list of video segments available at a headend.

The terminal then sends the request for the video segment to the headend for the headend to deliver (or re-deliver) the video segment from its start time. Upon receiving the request, equipment at the headend may re-compose a video sequence for the IPG screen with the time-shifted video segment, then encode and deliver the re-composed video sequence for the IPG screen to the viewer. The equipment at the headend may execute this process only for the requesting viewers.

V. MESSAGING PROTOCOL

Returning attention to the system 600 of Fig. 6, the following describes a messaging protocol for communicating between the major components of the system 600. The messaging protocol is described in relation to Figs. 8-11. Although a specific messaging protocol is described below, the present invention is not meant to be limited to that specific protocol.

First, the "source" transport stream generator (TSG) 604 communicates to a terminal 606 via, for example, a one-way in-band channel. The communication may be, for example, in the form of a "demand-cast index table" within a private section of the MPEG transport stream. Fig. 8 depicts an embodiment for the content of the demand-cast index table. The content may include: (a) a table version number (which increments when the table content changes); (b) a list of available demand-cast streams; (c) an internet protocol (IP) address for the source TSG; (d) a MUX channel number within the source TSG, and (e) a time of day and day of week.

Second, the terminal 606 communicates with the session manager (SM) 602 via, for example, a one-way out-of-band return path. The return path may be implemented, for example, using a user datagram protocol/internet protocol (UDP/IP) service to connect the terminal 606 to a network controller (NC) at the SM 602. Fig. 9

depicts one embodiment for the contents of the messages sent from the terminal 606 to the SM 602. The message content as shown includes: (a) a demand-cast stream identification; (b) the terminal's identification; (c) the IP address of the source TSG; (d) the MUX channel number within the source TSG; and (e) the message information itself.

5 The message information may indicate: (1) an acquisition of the demand-cast stream by the terminal 606; (2) a release of the demand-cast stream by the terminal 606; or (3) a request for the demand-cast stream by the terminal 606.

Third, the SM 602 communicates with the source TSG 604 via, for example, a two-way communications channel. The two-way communications channel
10 may comprise, for example, a TCP/IP connection over an Ethernet network. Fig. 10 depicts one embodiment for the contents of the messages sent from the SM 602 to the TSG 604. The message content as shown includes: (a) the demand-cast stream identification; (b) the MUX channel number within the source TSG; and (c) a message/command from a set of messages/commands. The set of messages/commands
15 include: (1) demand-cast stream released (no longer acquired by any terminal); (2) demand-cast stream requested; and (3) a reset command.

Messages from the SM 602 to the TSG 604 may be acknowledged by the TSG 604. Fig. 11 depicts one embodiment for the contents of the acknowledgement messages sent by the TSG 604 back to the SM 602. An acknowledgement message as
20 shown includes: (a) the demand-cast stream ID; (b) the MUX channel number; (c) the TSG's address; and (d) the acknowledgement itself. The acknowledgement may convey acknowledgement of: (1) release of the demand-cast stream; (2) request for the demand-cast stream; or (3) reset of the TSG 604.

VI. STREAM STATUS AND CONVERSIONS OF STATUS

25 The following relate to stream status and conversions of status in accordance with a preferred embodiment of the present invention. Although a specific embodiment of stream status and conversions of status is described below, the present invention is not meant to be limited to that specific embodiment.

1. Stream Status Within Multiplexed Transport Stream

The TSG 604 models bandwidth usage for each multiplexed transport stream that it is managing. Each demand-cast stream within each transport stream may be either active or inactive. Active streams are currently being multiplexed into the transport stream. Inactive streams are not currently being multiplexed into the transport stream.

Figure 12 depicts an example showing statuses of active demand-cast streams in a multiplexed transport stream generated by a TSG. For each demand-cast stream, TSG assigns status with respect to the streams intended multiplex. Demand-cast stream status, in context of the TSG, are:

1) 'active' streams are in the multiplex

1.1) 'acquired' demand-cast streams are in the multiplex and are used by at least one STT. They are referenced in the demand-cast index table in the private section of the transport stream. They are not candidates for removal.

1.2) 'released' demand-cast streams are in the multiplex and are not used by any STT. They are referenced in the demand-cast index table. They can become 'passive.'

1.2.1) 'passive' demand-cast streams are also technically 'released'.

They are in the multiplex and are not used by any STT. They are not referenced in the demand-cast index table. They are typically a small fraction of the 'released' demand-cast streams. The oldest few 'released' demand-cast streams are forced to 'inactive' status by a maintenance thread. They are candidates for removal.

2) 'inactive' demand-cast streams are not in the IPG multiplex. They are not referenced in the demand-cast index table. They may be inserted in the multiplex.

Note that the TSG may remove all the 'passive' demand-cast streams from their respective multiplexes and replace them with null packets. It is however

advantageous to leave 'passive' demand-cast streams in the multiplex because if a STT attempts to acquire it, latency will be minimized.

2. Conversions of Status

The TSG receives messages from the SM. Messages received from the SM are:

- 5 1) "request demand-cast stream"
- 2) "release demand-cast stream" The "release demand-cast stream" message includes the maximum number of STTs that were registered to the demand-cast stream.
- 3) "reset"

10 While specific embodiments and applications of the present invention have been illustrated and described, it is to be understood that the invention is not limited to the precise configuration and components disclosed above. Various modifications and variations which will be apparent to those skilled in the art may be made in the arrangement, operation, and details of the method and apparatus of the present invention
15 disclosed herein without departing from the spirit and scope of the invention.

 For example, in another embodiment, video segments in a small-size window may be overlaid on a channel program being broadcast. Such video segments may be set to begin playing at the time of viewer selection of the channel. In this embodiment, the video-segment overlay may be configured to be removable by a viewer
20 at any time by use of a remote button to turn off this overlay feature. Since the video segments may be stored at equipment at a headend, a small-size window displaying a video segment may be overlaid on any desired program channel at the headend and the resultant video sequence delivered to the viewer. In this embodiment, the overlay position on the screen may be configurable by a viewer. The viewer would be able to
25 define the coordinates of the overlay position and send the coordinates to the headend. Upon receiving the coordinates, the equipment and the headend would re-composite the video sequence with the video segment overlay at the defined coordinates.

In an additional embodiment, the short-time duration video segments may be delivered as part of a video-on-demand (VOD) service. Such a VOD service may be separate from an IPG context, but may be merged with a shopping mall context.

5 In yet another embodiment, the video segments may be delivered with a custom interactive program guide (custom IPG). Such a custom IPG may be delivered in a channel independent of the program channels and independent of the regular (non-custom) IPG channel. In one application, the video segments may be used to provide functionality to make offers for sale as part of a home shopping network integrated with the custom IPG.

10 FIG. 13 depicts a matrix 1300 of program guide data configured to present a different video for each PID. Matrix 1300 can be used to support, for example, look-ahead time selection in which a preview clip is provided for each PID. In this case, the guide data in the PIDs is the same (e.g., a list of eight channels) and the video data varies from PID to PID. Thus, each PID in matrix 1500 carries its own preview video clip for
15 its channel.

For matrix 1300, the guide data (represented as g1 in FIG. 13) can be encoded along with the first image of a reference PID as an I frame. Each of the remaining non-reference PIDs can be encoded independently as a different video sequence (e.g., a1, a2, a3, and so on). However, since the guide data portion (g1) is the
20 same for the PIDs, it can be omitted from processing and transmission.

Specifically, for time t1, the guide and video data for one of the PIDs (e.g., g1, v1 for PID1) can be encoded as the reference I frame. Subsequently, the video portions of the remaining pictures within the GOP for this PID can be encoded based on the reference I frame. The video portions at time index t1 for each of the remaining PIDs
25 (e.g., PID2 through PID8) can be encoded as an I frame. Alternatively, the video portion at time index t1 for each remaining PID can be coded as a P frame based on the reference I frame.

For example, the guide data (g1) and video data (v1) for PID1 at time index t1 can be encoded as the reference I frame. For the next picture of PID1 at time
30 index t2, the video data (v2) is extracted and encoded as a B picture based, in part, on the video data (v1) at time index t1. The guide data (g1) at time t2 can be omitted from

processing. The encoding for PID1 continues in similar manner for the remaining pictures at time indices t3 through t15. For PID2, the video data (a1) at time index t1 can be coded as an I frame, and the video data (a2, a3, and so on) for time indices t2 through t15 can be encoded as P and B frames based on the I frame generated for PID2 at time index t1.

The decoding for data structure 1300 can be performed (e.g., at the STT) as follows. Initially, the reference I frame is constructed and stored. If a particular PID is selected for viewing, the video sequence for that PID is constructed and combined with the previously constructed and stored guide data. The decoded video sequence is thus presented along with the guide data available in the decoded reference frame.